

Model Accuracy and Runtime Tradeoff in Distributed Deep Learning: A Systematic Study

Suyog Gupta^{1,*}, Wei Zhang^{1,*}, Fei Wang²

¹IBM T. J. Watson Research Center, Yorktown Heights, NY. *Equal Contribution.

²Department of Healthcare Policy and Research, Weill Cornell Medical College, New York City, NY

Abstract—Deep learning with a large number of parameters requires distributed training, where model accuracy and runtime are two important factors to be considered. However, there has been no systematic study of the tradeoff between these two factors during the model training process. This paper presents Rudra, a parameter server based distributed computing framework tuned for training large-scale deep neural networks. Using variants of the asynchronous stochastic gradient descent algorithm we study the impact of synchronization protocol, stale gradient updates, minibatch size, learning rates, and number of learners on runtime performance and model accuracy. We introduce a new learning rate modulation strategy to counter the effect of stale gradients and propose a new synchronization protocol that can effectively bound the staleness in gradients, improve runtime performance and achieve good model accuracy. Our empirical investigation reveals a principled approach for distributed training of neural networks: the mini-batch size per learner should be reduced as more learners are added to the system to preserve the model accuracy. We validate this approach using commonly-used image classification benchmarks: CIFAR10 and ImageNet.

I. INTRODUCTION

Deep neural network based models have achieved unparalleled accuracy in cognitive tasks such as speech recognition, object detection, and natural language processing [18]. For certain image classification benchmarks, deep neural networks have been touted to even surpass human-level performance [13, 11]. Such accomplishments are made possible by the ability to perform fast, supervised training of complex neural network architectures using large quantities of labeled data. Training a deep neural network translates into solving a non-convex optimization problem in a very high dimensional space, and in the absence of a solid theoretical framework to solve such problems, practitioners are forced to rely on trial-and-error empirical observations to design heuristics that help obtain a well-trained model[2]. Naturally, fast training of deep neural network models can enable rapid evaluation of different network architectures and facilitate a more thorough hyper-parameter optimization for these models. Recent years have seen a resurgence of interest in deploying large-scale computing infrastructure designed specifically for training deep neural networks. Some notable efforts in this direction include distributed computing infrastructure using thousands of CPU cores [4, 7], high-end graphics processors (GPUs)[16], or a combination of CPUs and GPUs [5].

The large-scale deep learning problem can hence be viewed as a confluence of elements from machine learning (ML) and high-performance computing (HPC). Much of the work in the ML community is focused on non-convex optimization, model

selection, and hyper-parameter tuning to improve the neural network's performance (measured as classification accuracy) while working under the constraints of the computational resources available in a single computing node (CPU with or without GPU acceleration). From a HPC perspective, prior work has addressed, to some extent, the problem of accelerating the neural network training by mapping the asynchronous version of mini-batch stochastic gradient descent (SGD) algorithm onto multiple computing nodes. Contrary to the popular belief that asynchrony necessarily improves model accuracy, we find that adopting the approach of scale-out deep learning using asynchronous-SGD, gives rise to complex interdependencies between the training algorithm's hyperparameters and the distributed implementation's design choices (synchronization protocol, number of learners), ultimately impacting the neural network's accuracy and the overall system's runtime performance.

In this paper we present *Rudra*, a parameter server based deep learning framework to study these interdependencies and undertake an empirical evaluation with public image classification benchmarks. Our key contributions are:

- 1) A systematic technique (vector clock) for quantifying the staleness of gradient descent parameter updates.
- 2) An investigation of the impact of the interdependence of training algorithm's hyperparameters (mini-batch size, learning rate (gradient descent step size)) and distributed implementation's parameters (gradient staleness, number of learners) on the neural network's classification accuracy and training time.
- 3) A new learning rate tuning strategy that reduces the effect of stale parameter updates.
- 4) A new synchronization protocol to reduce network bandwidth overheads while achieving good classification accuracy and runtime performance.
- 5) An observation that to maintain a given level of model accuracy, it is necessary to reduce the mini-batch size as the number of learners is increased. This suggests a hard limit on the amount of parallelism that can be exploited in training a given model.

II. BACKGROUND

A neural network computes a parametric, non-linear transformation $f_{\theta} : X \mapsto Y$, where θ represents a set of adjustable parameters (or weights). In a supervised learning context (such as image classification), X is the input image and Y corresponds to the label assigned to the image. A deep neural

network organizes the parameters θ into multiple *layers*, each of which consists of a linear transformation followed by a non-linear function such as sigmoid, tanh, etc. In a feed-forward deep neural network, the layers are arranged hierarchically such that the output of the layer $l - 1$ feeds into the input of layer l . The terminal layer generates the network's output $\hat{Y} = f_\theta(X)$, corresponding to the input X .

A neural network training algorithm seeks to find a set of parameters θ^* that minimizes the discrepancy between \hat{Y} and the ground truth Y . This is usually accomplished by defining a differentiable cost function $C(\hat{Y}, Y)$ and iteratively updating each of the model parameters using some variant of the gradient descent algorithm:

$$E_m = \frac{1}{m} \sum_{s=1}^m C(\hat{Y}_s, Y_s), \quad (1a)$$

$$\nabla \theta^{(k)}(t) = \left(\partial E_m / \partial \theta^{(k)} \right)(t), \quad (1b)$$

$$\theta^{(k)}(t+1) = \theta^{(k)}(t) - \alpha(t) \nabla \theta^{(k)}(t) \quad (1c)$$

where $\theta^{(k)}(t)$ represents the k^{th} parameter at iteration t , α is the step size (also known as the learning rate) and m is the batch size. The batch gradient descent algorithm sets m to be equal to the total number of training examples N . Due to the large amount of training data, deep neural networks are typically trained using the Stochastic Gradient Descent (SGD), where the parameters are updated with a randomly selected training example (X_s, Y_s) . The performance of SGD can be improved by computing the gradients using a *mini-batch* containing $m = \mu \ll N$ training examples.

Deep neural networks are generally considered hard to train [2, 10, 23] and the trained model's generalization error depends strongly on hyperparameters such as the initializations, learning rates, mini-batch size, network architecture, etc. In addition, neural networks are prone to overfit the data. Regularization methods (e.g., weight decay and dropout) [16] applied during training have been shown to combat overfitting and reduce the generalization error.

Scale-out deep learning: A typical implementation of distributed training of deep neural networks consists of a master (parameter server) that orchestrates the work among one or more slaves (learners). Each learner does the followings:

- 1) `getMinibatch`: Select randomly a mini-batch of examples from the training data.
- 2) `pullWeights`: Request the parameter server for the current set of weights/parameters.
- 3) `calcGradient`: Compute gradients based on the training error for the current mini-batch (equation 1b).
- 4) `pushGradient`: Send the computed gradients to the parameter server

The parameter server maintains a global view of the model weights and performs the following functions:

- 1) `sumGradients`: Receive and accumulate the gradients from the learners.
- 2) `applyUpdate`: Multiply the accumulated gradient by the learning rate and update the weights (equation 1c)

Learners exploit *data parallelism* by each maintaining a copy of the entire model, and training independently over a

unique mini-batch. The *model parallelism* approach augments this framework by splitting the neural network model across multiple learners. With model parallelism, each learner trains only a portion of the network; edges that cross learner boundaries must be synchronized before gradients can be computed for the entire model.

Several different synchronization strategies are possible. The most commonly used one is the asynchronous protocol, in which the learners work completely independently of each other and the parameter server. Section III will discuss three different synchronization strategies, each associated with a unique tradeoff between model accuracy and runtime.

III. DESIGN AND IMPLEMENTATION

A. Terminology

Throughout the paper, we use the following definitions:

- **Parameter Server**: a server that holds the model weights. [22] describes a typical parameter server using a distributed key-value store to synchronize state between processes. The parameter server collects gradients from learners and updates the weights accordingly.
- **Learner**: A computing process that can calculate weight updates (gradients).
- μ : mini-batch size.
- α : learning rate.
- λ : number of learners.
- **Epoch**: a pass through the entire training dataset.
- **Timestamp**: we use a scalar clock [20] to represent weights timestamp ts_i , starting from $i = 0$. Each weight update increments the timestamp by 1. The timestamp of a gradient is the same as the timestamp of the weight used to compute the gradient.
- σ : staleness of the gradient. A gradient with timestamp ts_i is pushed to the parameter server with current weight timestamp ts_j , where $ts_j \geq ts_i$. We define the staleness of this gradient σ as $j - i$.
- $\langle \sigma \rangle$, average staleness of gradients. The timestamps of the set of n gradients that triggers the advancement of weights timestamp from ts_{i-1} to ts_i form a vector clock [17] $\langle ts_{i_1}, ts_{i_2}, \dots, ts_{i_n} \rangle$, where $\max\{i_1, i_2, \dots, i_n\} < i$. The average staleness of gradients $\langle \sigma \rangle$ is defined as:

$$\langle \sigma \rangle = (i - 1) - \text{mean}(i_1, i_2, \dots, i_n) \quad (2)$$

- **Hardsync protocol**: To advance weights timestamp from ts_i to ts_{i+1} , each learner calculates exactly one mini-batch and sends its gradient $\nabla \theta_l$ to the parameter server. The parameter server averages the gradients and updates the weights according to Equation (3), then broadcasts the new weights to all learners. Staleness in the hardsync protocol is always zero.

$$\begin{aligned} \nabla \theta^{(k)}(i) &= \frac{1}{\lambda} \sum_{l=1}^{\lambda} \nabla \theta_l^{(k)} \\ \theta^{(k)}(i+1) &= \theta^{(k)}(i) - \alpha \nabla \theta^{(k)}(i) \end{aligned} \quad (3)$$

- **Async protocol**: Each learner calculates the gradients and asynchronously pushes/pulls the gradients/weights




Fig. 1. Rudra-base architecture

to/from parameter server. The Async weight update rule is given by:

$$\begin{aligned} \nabla \theta^{(k)}(i) &= \nabla \theta_l^{(k)}, L_l \in L_1, \dots, L_\lambda \\ \theta^{(k)}(i+1) &= \theta^{(k)}(i) - \alpha \nabla \theta^{(k)}(i) \end{aligned} \quad (4)$$

Gradient staleness may be hard to control due to the asynchrony in the system. [7] describe *Downpour SGD*, an implementation of the Async protocol for a commodity scale-out system in which the staleness can be as large as hundreds.

- *n*-softsync protocol: Each learner pulls the weights from the parameter server, calculates the gradients and pushes the gradients to the parameter server. The parameter server updates the weights after collecting at least $c = \lfloor (\lambda/n) \rfloor$ gradients. The splitting parameter n can vary from 1 to λ . The *n*-softsync weight update rule is given by:

$$\begin{aligned} c &= \lfloor (\lambda/n) \rfloor \\ \nabla \theta^{(k)}(i) &= \frac{1}{c} \sum_{l=1}^c \nabla \theta_l^{(k)}, L_j \in L_1, \dots, L_\lambda \\ \theta^{(k)}(i+1) &= \theta^{(k)}(i) - \alpha \nabla \theta^{(k)}(i) \end{aligned} \quad (5)$$

In Section V-A we will show that in a homogeneous cluster where each learner proceeds at roughly the same speed, the staleness of the model can be empirically bounded at $2n$. Note that when n is equal to λ , the weight update rule at the parameter server is exactly the same as in Async protocol.

B. Rudra-base System Architecture

Figure 1 illustrates the parameter server design that we use to study the interplay of hyperparameter tuning and system scale-out factor. This system implements both hardsync and *n*-softsync protocols. The arrows between each entity represent a (group of) MPI message(s), except the communication between *Learner* and *Data Server*, which is achieved by a global file system. We describe each entity's role and its implementation below.

Learner is a single-process multithreaded SGD solver. Before training each mini-batch, a learner pulls the weights and the corresponding timestamp from the parameter server. A learner reduces the pullWeights traffic by first inquiring the timestamp from the parameter server: if the timestamp is as old as the local weights', then this learner does not pull the weights. After training the mini-batch, learner sends the

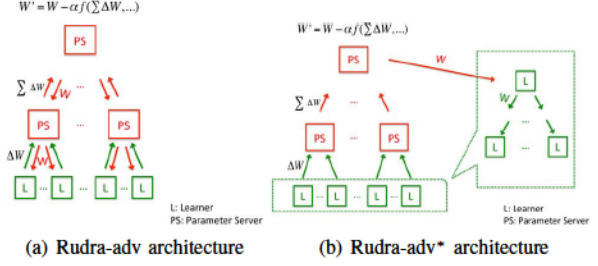


Fig. 2. Rudra-adv architecture

gradients along with gradients' timestamp to parameter server. The size of pull and push messages is the same as the model size plus the size of scalar timestamp equal to one.

Data Server is hosted on IBM GPFS, a global file system. Each learner has an I/O thread, which prefetches the mini-batch via random sampling prior to training. Prefetching is completely overlapped with the computation.

Parameter Server is a multithreaded process, that accumulates gradients from each learner and applies update rules according to Equations (3–5). In this study, we implemented hardsync protocol and *n*-softsync protocol. Learning rate is configured differently in either protocol. In hardsync protocol, the learning rate is multiplied by a factor $\sqrt{\lambda\mu/B}$, where B is the batch size of the reference model. In the *n*-softsync protocol, the learning rate is multiplied by the reciprocal of staleness. We demonstrate in Section V-A that this treatment of learning rate in *n*-softsync can significantly improve the model accuracy. Parameter server records the vector clock of each weight update to keep track of the the average staleness. When a specified number of epochs are trained, parameter server shuts down each learner.

Statistics Server is a multithreaded process that receives the training error from each learner and receives the model from the parameter server at the end of each epoch and tests the model. It monitors the model training quality.

This architecture is non-blocking everywhere except for pushing up gradients and pushing down weights, which are blocking MPI calls (e.g. `MPI_Send`). Parameter server handles each incoming message one by one (the message handling itself is multithreaded). In this way, we can precisely control how each learner's gradients are received and handled by the parameter server. The purpose of Rudra-base is to control the noise of the system, so that we can study the interplay of scale-out factor and the hyperparameter selection. For a moderately-sized dataset like CIFAR-10, Rudra-base shows good scale-out factor (see Section V-B).

C. Rudra-adv and Rudra-adv* System Architecture

To achieve high classification accuracy, the required model size may be quite large (e.g. hundreds of MBs). In many cases, to achieve best possible model accuracy, mini-batch size μ must be small, as we will demonstrate in Section V-B. In order to meet these requirements with acceptable performance, we implemented Rudra-adv and Rudra-adv*.

Rudra-adv system architecture. Rudra-base clearly is not a scalable solution when the model gets large. Under ideal circumstances (see Section IV-A for our experimental hardware system specification), a single learner pushing a model of 300 MB (size of a typical deep neural network, see section IV-B) would take more than 10 ms to transfer this data. If 16 tasks are sending 300 MB to the same receiver and there is link contention, it would take over a second for the messages to be delivered.

To alleviate the network traffic to parameter server, we build a parameter server group that forms a tree. We co-locate each tree leaf node on the same node as the learners for which it is responsible. Each node in the parameter server group is responsible for averaging the gradients sent from its learners and relaying the averaged gradient to its parent. The root node in the parameter server group is responsible for applying weight update and broadcasting the updated weights. Each non-leaf node pulls the weights from its parent and responds to its children’s weight pulling requests. Rudra-adv can significantly improve performance compared to Rudra-base and manage to scale out to large model and small μ , while maintaining the control of the gradients’ staleness. Figure 2(a) illustrates the system architecture for Rudra-adv. Red boxes represent the parameter server group, in which the gradients are pushed and aggregated upwards. Green boxes represent learners, each learner pushes the gradient to its parameter server parent and receives weights from its parameter server parent. The key difference between Rudra-adv and a sharded parameter server design (e.g., Distbelief [7] and Adam [4]) is that the weights maintained in Rudra-adv have the same timestamps whereas shared parameter servers maintain the weights with different timestamps. Having consistent weights makes the analysis of hyperparameter/scale-out interplay much more tractable.

Rudra-adv* system architecture. We built Rudra-adv* to further improve the runtime performance in two ways:

Broadcast weights within learners. To further reduce the traffic to the parameter server group, we form a tree within all learners and broadcast the weights down this tree. In this way the network links to/from learners are also utilized.

Asynchronous pushGradient and pullWeights. Ideally, one would use MPI non-blocking send calls to asynchronously send gradients and weights. However, depending on the MPI implementation, it is difficult to guarantee if the non-blocking send calls make progress in the background [1]. Therefore we open additional communication threads and use MPI blocking send calls in the threads. Each learner process runs two additional communication threads: the pullWeights thread and pushGradient thread. In this manner, computing can continue without waiting for the communication. Note that since we need to control μ (the smaller μ is, the better model converges, as we demonstrate in Section V-B), we must guarantee that the learner pushes each calculated gradient to the server. Alternatively, one could locally accrue gradients and send the sum, as in [7], however that will effectively increase μ . For this

Implementation	Communication overlap (%)
Rudra-base	11.52
Rudra-adv	56.75
Rudra-adv*	99.56

TABLE I
COMMUNICATION OVERLAP MEASURED IN RUDRA-BASE, RUDRA-ADV, RUDRA-ADV* FOR AN ADVERSARIAL SCENARIO, WHERE THE MINI-BATCH SIZE IS THE SMALLEST POSSIBLE FOR 4-WAY MULTI-THREADED LEARNERS, MODEL SIZE 300MB, AND THERE ARE ABOUT 60 LEARNERS.

reason, the pushGradient thread cannot start sending the current gradient before the previous one has been delivered. As demonstrated in Table I that as long as we can optimize the use of network links, this constraint has no bearing on the runtime performance, even when μ is extremely small. In contrast, pullWeights thread has no such constraint – we maintain a computation buffer and a communication buffer for pullWeights thread, and the communication always happens in the background. To use the newly received weights only requires a pointer swap. Figure 2(b) illustrates the system architecture for Rudra-adv*. Different from Rudra-adv, each learner continuously receives weights from the weights downcast tree, which consists of the top level parameter server node and all the learners.

We measure the communication overlap by calculating the ratio between computation time and the sum of computation and communication time. Table I records the the communication overlap for Rudra-base, Rudra-adv, and Rudra-adv* in an adversarial scenario. Rudra-adv* can almost completely overlap computation with communication. Rudra-adv* can scale out to very large model size and work with smallest possible size of mini-batch. In Section V-E, we demonstrate Rudra-adv*’s effectiveness in improving runtime performance while achieving good model accuracy.

IV. METHODOLOGY

A. Hardware and software environment

We deploy the Rudra distributed deep learning framework on a P775 supercomputer. Each node of this system contains four eight-core 3.84 GHz POWER7 processors, one optical connect controller chip and 128 GB of memory. A single node has a theoretical floating point peak performance of 982 Gflop/s, memory bandwidth of 512 GB/s and bi-directional interconnect bandwidth of 192 GB/s.

The cluster operating system is Red Hat Enterprise Linux 6.4. To compile and run Rudra we used the IBM xLC compiler version 12.1 with the `-O3 -q64 -qsmp` options, ESSL for BLAS subroutines, and IBM MPI (IBM Parallel Operating Environment 1.2).

B. Benchmark datasets and neural network architectures

To evaluate Rudra’s scale-out performance we employ two commonly used image classification benchmark datasets: CIFAR10 [15] and ImageNet [9]. The CIFAR10 dataset comprises of a total of 60,000 RGB images of size 32×32 pixels partitioned into the training set (50,000 images) and the test set (10,000 images). Each image belongs to one of the

10 classes, with 6000 images per class. For this dataset, we construct a deep convolutional neural network (CNN) with 3 convolutional layers each followed by a subsampling/pooling layer. The output of the 3rd pooling layer connects, via a fully-connected layer, to a 10-way softmax output layer that generates a probability distribution over the 10 output classes. This neural network architecture closely mimics the CIFAR10 model (cifar10_full.prototxt) available as a part of the open-source Caffe deep learning package [14]. The total number of trainable parameters in this network are ~ 90 K, resulting in the model size of ~ 350 kB when using 32-bit floating point data representation. The neural network is trained using momentum-accelerated mini-batch SGD with a batch size of 128 and momentum set to 0.9. As a data preprocessing step, the per-pixel mean is computed over the entire training dataset and subtracted from the input to the neural network. The training is performed for 140 epochs and results in a model that achieves 17.9% misclassification error rate on the test dataset. The base learning rate α_0 is set to 0.001 and reduced by a factor of 10 after the 120th and 130th epoch. This learning rate schedule proves to be quite essential in obtaining the low test error of 17.9%.

Our second benchmark dataset is collection of natural images used as a part of the 2012 edition of the ImageNet Large Scale Visual Recognition Challenge (ILSVRC 2012). The training set is a subset of the hand-labeled ImageNet database and contains 1.2 million images. The validation dataset has 50,000 images. Each image maps to one of the 1000 non-overlapping object categories. The images are converted to a fixed resolution of 256×256 to be used input to a deep convolution neural network. For this dataset, we consider the neural network architecture introduced in [16] consisting of 5 convolutional layers and 3 fully-connected layers. The last layer outputs the probability distribution over the 1000 object categories. In all, the neural network has ~ 72 million trainable parameters and the total model size is 289 MB. The network is trained using momentum-accelerated SGD with a batch size of 256 and momentum set to 0.9. Similar to the CIFAR10 benchmark, per-pixel mean computed over the entire training dataset is subtracted from the input image feeding into the neural network. To prevent overfitting, a weight regularization penalty of 0.0005 is applied to all the layers in the network and a dropout of 50% is applied to the 1st and 2nd fully-connected layers. The initial learning rate α_0 is set equal to 0.01 and reduced by a factor of 10 after the 15th and 25th epoch. Training for 30 epochs results in a top-1 error of 43.95% and top-5¹ error of 20.55% on the validation set.

V. EVALUATION

In this section we present results of evaluation of our scale-out deep learning training implementation. For an initial design

¹The top-5 error corresponds to the fraction of samples where the correct label does not appear in the top-5 labels considered most probable by the model




Fig. 3. Average staleness $\langle \sigma \rangle$ of the gradients as a function of the weight update step at the parameter server when using (a) 1-softsync, 2-softsync and (b) λ -softsync protocol. Inset in (b) shows the distribution of the gradient staleness values for λ -softsync protocol. Number of learners λ is set to 30.

space exploration, we use the CIFAR10 dataset and Rudra-base system architecture. Subsequently we extend our findings to the ImageNet dataset using the Rudra-adv and Rudra-adv* system architectures.

A. Stale gradients

In the hardsync protocol introduced in section III-A, the transition from $\theta(i)$ to $\theta(i+1)$ involves aggregating the gradients calculated with $\theta(i)$. As a result, each of the gradients $\nabla \theta_i$ carries with it a staleness σ equal to 0. However, departing from the hardsync protocol towards either the n -softsync or the Async protocol inevitably adds staleness to the gradients, as a subset of the learners contribute gradients calculated using weights with timestamp earlier than the current timestamp of the weights at the parameter server.

To measure the effect of gradient staleness when using the n -softsync protocol, we use the CIFAR10 dataset and train the neural network described in section IV-B using $\lambda = 30$ learners. For the 1-softsync protocol, the parameter server updates the current set of weights when it has received a total of 30 gradients from the learners. Similarly, the 2-softsync protocol forces the parameter server to accumulate $\lambda/2 = 15$ gradient contributions from the learners before updating the weights. As shown in Figure 3(a) the average staleness $\langle \sigma \rangle$ for the 1-softsync and 2-softsync protocols remains close to 1 and 2, respectively. In the 1-softsync protocol, the staleness σ_{L_i} for the gradients computed by the learner L_i takes values 0, 1, or 2, whereas $\sigma_{L_i} \in \{0, 1, 2, 3, 4\}$ for the 2-softsync protocol. Figure 3(b) shows the gradient staleness for the n -softsync protocol where $n = \lambda = 30$. In this case, the parameter server updates the weights after receiving a gradient from any of the learners. A large fraction of the gradients have staleness close to 30, and only with a very low probability (< 0.0001) does σ exceed $2n = 60$. These measurements show that, in general, $\sigma_{L_i} \in \{0, 1, \dots, 2n\}$ and $\langle \sigma \rangle = n$ for our implementation of the n -softsync protocol.

Modifying the learning rate for stale gradients: In our experiments with the n -softsync protocol we found it beneficial, and at times necessary, to modulate the learning rate α to take into account the staleness of the gradients. For the n -softsync protocol, we set the learning rate as:

$$\alpha = \alpha_0 / \langle \sigma \rangle = \alpha_0 / n \quad (6)$$

where α_0 is the learning rate used for the baseline (control) experiment: $\mu = 128$, $\lambda = 1$. Figure 4 shows a set of




Fig. 4. *Effect of learning rate modulation strategy*: Dividing the learning rate by the average staleness aids in better convergence and achieves lower test error when using the n -softsync protocol. Number of learners, $\lambda = 30$; mini-batch size $\mu = 128$.




Fig. 5. (σ, μ, λ) tradeoff curves for the hardsync protocol. The dashed black line represents the 17.9% test error achieved by the baseline model $(\sigma, \mu, \lambda) = (0, 128, 1)$ on the CIFAR10 dataset.

representative results illustrating the benefits of adopting this learning rate modulation strategy. We show the evolution of the test error on the CIFAR10 dataset as a function of the training epoch for two different configurations of the n -softsync protocol ($n = 4, n = 30$) and set the number of learners, $\lambda = 30$. In both these configurations, setting the learning rate in accordance with equation (6) results in lower test error as compared with the cases where the learning rate is set to α_0 . Surprisingly, the configuration 30-softsync, $\lambda = 30$, $\alpha = \alpha_0$ fails to converge and shows a constant high error rate of 90%. Reducing the learning rate by a factor $\langle \sigma \rangle = n = 30$ makes the model with much lower test error².

B. (σ, μ, λ) tradeoff curves

Hyperparameter optimization plays a central role in obtaining good accuracy from neural network models [3]. For the SGD training algorithm, this includes a search over the neural network’s training parameters such as learning rates, weight regularization, depth of the network, mini-batch size etc. in order to improve the quality of the trained neural network model (quantified as the error on the validation dataset). Additionally, when distributing the training problem across multiple learners, parameters such as the number of learners and the synchronization protocol enforced amongst the learners impact not only the runtime of the algorithm but also the quality of the trained model.

An exhaustive search over the space defined by these parameters for joint optimization of the runtime performance and the model quality can prove to be a daunting task even

²Although not explored as a part of this work, it is certainly possible to implement a finer-grained learning rate modulation strategy that depends on the staleness of each of gradients computed by the learners instead of the average staleness. Such a strategy should apply smaller learning rates to staler gradients

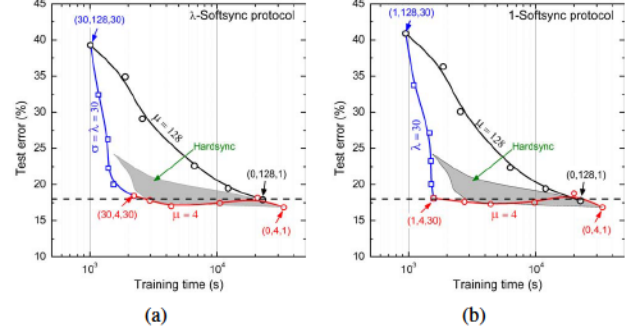


Fig. 6. (σ, μ, λ) tradeoff curves for (a) λ -softsync protocol and (b) 1-softsync protocol. Shaded region in shows the region bounded by $\mu = 128$, $\lambda = 30$, and $\mu = 4$ contours for the hardsync protocol. $\lambda \in \{1, 2, 4, 10, 18, 30\}$ and $\mu \in \{4, 8, 16, 32, 64, 128\}$. Note that for $\lambda = 1$, n -softsync protocol degenerates to the hardsync protocol

for a small model such as that used for the CIFAR10 dataset, and clearly intractable for models and datasets the scale of ImageNet. To develop a better understanding of the interdependence among the various tunable parameters in the distributed deep learning problem, we introduce the notion of (σ, μ, λ) tradeoff curves. A (σ, μ, λ) tradeoff curve plots the error on the validation set (or test set) and the total time to train the model (wall clock time) for different configurations of average gradient staleness $\langle \sigma \rangle$, mini-batch size per learner μ , and the number of learners λ . The configurations (σ, μ, λ) that achieve the virtuous combination of low test error and small training time are preferred and form ideal candidates for further hyperparameter optimization.

We run³ the CIFAR10 benchmark for $\lambda \in \{1, 2, 4, 10, 18, 30\}$ and $\mu \in \{4, 8, 16, 32, 64, 128\}$. Figure 5 shows a set of (σ, μ, λ) curves for the hardsync protocol i.e. $\sigma = 0$. The baseline configuration with $\lambda = 1$ learner and mini-batch size $\mu = 128$ achieves a test error of 17.9%. With the exception of modifying the learning rate as $\alpha = \alpha_0 \sqrt{\mu \lambda / 128}$, all the other neural network’s hyperparameters were kept unchanged from the baseline configuration while generating the data points for different values of μ and λ . Note that it is possible to achieve test error lower than the baseline by reducing the mini-batch size from 128 to 4. However, this configuration (indicated on the plot as $(\sigma, \mu, \lambda) = (0, 4, 1)$) increases training time compared with the baseline. This is primarily due to the fact that the dominant computation performed by the learners involves multiple calls to matrix multiplication (GEMM) to compute WX where samples in a mini-batch form columns of the matrix X . Reducing the mini-batch size cause a proportionate decrease in the GEMM throughput and slower processing of the mini-batch by the learner.

In Figure 5, the contour labeled $\mu = 128$ is the configurations with the mini-batch size per learner is kept constant at 128 and the number of learners is varied from $\lambda = 1$ to $\lambda = 30$.

³The mapping between λ and the number of computing nodes η is $(\lambda, \eta) = \{(1, 1), (2, 1), (4, 2), (10, 4), (18, 4), (30, 8)\}$

The training time reduces monotonically with λ , albeit at the expense of an increase in the test error. Traversing along the $\lambda = 30$ contour from configuration $(\sigma, \mu, \lambda) = (0, 128, 30)$ to $(\sigma, \mu, \lambda) = (0, 4, 30)$ (i.e. reducing the mini-batch size from 128 to 4) helps restore much of this degradation in the test error by partially sacrificing the speed-up obtained by the virtue of scaling out to 30 learners.




Fig. 7. Speed-up in the training time for mini-batch size and (a) $\mu = 128$ (b) $\mu = 4$ for 3 different protocols: hardsync, λ -softsync, and 1-softsync. Speed-up numbers in (a) and (b) are calculated relative to $(\sigma, \mu, \lambda) = (0, 128, 1)$ and $(\sigma, \mu, \lambda) = (0, 4, 1)$, respectively.

Figure 6(a) shows (σ, μ, λ) tradeoff curves for the λ -softsync protocol. In this protocol, the parameter server updates the weights as soon as it receives a gradient from any of the learners. Therefore, as shown in section V-A the average gradient staleness $\langle \sigma \rangle = \lambda$ and $\sigma_{max} \leq 2\lambda$ with high probability. The learning rate is set in accordance with equation 6. All the other hyperparameters are left unchanged from the baseline configuration. Qualitatively, the (σ, μ, λ) tradeoff curves for λ -softsync look similar to those observed for the hardsync protocol. In this case, however, the degradation in the test error relative to the baseline for the $(\sigma, \mu, \lambda) = (30, 128, 30)$ configuration is much more pronounced. As observed previously, this increase in the test error can largely be mitigated by reducing the size of mini-batch processed by each learner ($\lambda = 30$ contour). Note, however, the sharp increase in the training time for the configuration $(\sigma, \mu, \lambda) = (30, 4, 30)$ as compared with $(\sigma, \mu, \lambda) = (30, 128, 30)$. The smaller mini-batch size not only reduces the computational throughput at each learner, but also increases the frequency of `pushGradient` and `pullWeights` requests at the parameter server. In addition, small mini-batch size increases the frequency of weight updates at the parameter server. Since in the Rudra-base architecture (section III-B), the learner does not proceed with the computation on the next mini-batch till it has received the updated gradients, the traffic at the parameter server and the more frequent weight updates can delay servicing the learner's `pullWeights` request, potentially stalling the gradient computation at the learner. Interestingly, all the configurations along the $\mu = 4$ contour show similar, if not better, test error as the baseline. For these configurations, the average staleness varies between 2 and 30. *From this empirical observation, we infer that a small mini-batch size per learner confers upon the training algorithm a fairly high degree of immunity to stale gradients.*

The 1-softsync protocol shows (σ, μ, λ) tradeoff curves (Figure 6(b)) that appear nearly identical to those observed

for the λ -softsync protocol. In this case, the average staleness is 1 irrespective of the number of learners. Since the parameter server waits for λ gradients to arrive before updating the weights, there is a net reduction in the `pullWeights` traffic at the parameter server (see section III-B). As a result, the 1-softsync protocol avoids the degradation in runtime observed in the λ -softsync protocol for the configuration with $\mu = 4$ and $\lambda = 30$. The distinction in terms of the runtime performance between these two protocols becomes obvious when comparing the speed-ups obtained for different mini-batch sizes (Figure 7). For $\mu = 128$, the 1-softsync and λ -softsync protocol demonstrate similar speed-ups over the baseline configuration for upto $\lambda = 30$ learners. In this case, the communication between the learners and the parameter server is sporadic enough to prevent the learners from waiting on the parameter server for updated weights. However, as the number of learners is increased beyond 30, the bottlenecks at the parameter server are expected to diminish the speed-up obtainable using the λ -softsync protocol. The effect of frequent `pushGradient` and `pullWeights` requests due to smaller at the parameter manifest clearly as the mini-batch size is reduced to 4, in which case, the λ -softsync protocol shows subdued speed-up compared with 1-softsync protocol. In either scenario, the hardsync protocol fares the worst in terms of runtime performance improvement when scaling out to large number of learners. The upside of adopting the hardsync protocol, however, is that it achieves substantially lower test error, even for large mini-batch sizes.

C. $\mu\lambda = \text{constant}$

In the hardsync protocol, given a configuration with λ learners and mini-batch size μ per learner, the parameter server averages the λ number of gradients reported to it by the learners. Using equations (1) and (3):

$$\begin{aligned} \nabla \theta^{(k)}(i) &= \frac{1}{\lambda} \sum_{l=1}^{\lambda} \nabla \theta_l^{(k)} = \frac{1}{\lambda} \sum_{l=1}^{\lambda} \left(\frac{1}{\mu} \sum_{s=1}^{\mu} \frac{\partial C(\hat{Y}_s, Y_s)}{\partial \theta_l^{(k)}(i)} \right) \\ &= \frac{1}{\mu\lambda} \sum_{s=1}^{\mu\lambda} \frac{\partial C(\hat{Y}_s, Y_s)}{\partial \theta^{(k)}(i)} \end{aligned} \quad (7)$$

The last step equation (7) is valid since each training example (X_s, Y_s) is drawn independently from the training set and also because the hardsync protocol ensures that all the learners compute gradients on identical set of weights i.e. $\theta_l^{(k)}(i) = \theta^{(k)}(i) \forall l \in \{1, 2, \dots, \lambda\}$. According to equation (7), the configurations $(\sigma, \mu, \lambda) = (0, \mu_0 \lambda_0, 1)$ and $(\sigma, \mu, \lambda) = (0, \mu_0, \lambda_0)$ are equivalent from the perspective of stochastic gradient descent optimization. In general, hardsync configurations with the same $\mu\lambda$ product are expected to give nearly⁴ the same test error.

In Table II we report the test error at the end of 140 epochs for configurations with $\mu\lambda = \text{constant}$. Interestingly, we find that even for the n -softsync protocol, configurations

⁴small differences in the final test error achieved may arise due to the inherent nondeterminism of random sampling in stochastic gradient descent and the random initialization of the weights.

	σ	μ	λ	Test error	Training time(s)
$\mu\lambda \approx 128$	1	4	30	18.09%	1573
	30	4	30	18.41%	2073
	18	8	18	18.92%	2488
	10	16	10	18.79%	3396
	4	32	4	18.82%	7776
	2	64	2	17.96%	13449
tiny $\mu\lambda \approx 256$	1	8	30	20.04%	1478
	30	8	30	19.65%	1509
	18	16	18	20.33%	2938
	10	32	10	20.82%	3518
	4	64	4	20.70%	6631
	2	128	2	19.52%	11797
	1	128	2	19.59%	11924
$\mu\lambda \approx 512$	1	16	30	23.25%	1469
	30	16	30	22.14%	1502
	18	32	18	23.63%	2255
	10	64	10	24.08%	2683
	4	128	4	23.01%	7089
$\mu\lambda \approx 1024$	1	32	30	27.16%	1299
	30	32	30	27.27%	1420
	18	64	18	28.31%	1713
	1	128	10	29.83%	2551
	10	128	10	29.90%	2626

TABLE II
Results on CIFAR10 benchmark: TEST ERROR AT THE END OF 140 EPOCHS AND TRAINING TIME FOR (σ, μ, λ) CONFIGURATIONS WITH $\mu\lambda = \text{CONSTANT}$.

that maintain $\mu\lambda = \text{constant}$ achieve comparable test errors. At the same time, the test error turns out to be rather independent of the staleness in the gradients for a given $\mu\lambda$ product. For instance, Table II shows that when $\mu\lambda \approx 128$, but the (average) gradient staleness is allowed to vary between 1 and 30, the test error stays ~ 18 -19%. Although this result may seem counter-intuitive, a plausible explanation emerges when considering the measurements shown earlier in Figure 3, that our implementation of the n -softsync protocol achieves an average gradient staleness of n while bounding the maximum staleness at $2n$. Consequently, at any stage in the gradient descent algorithm, the weights being used by the different learners $(\theta_i(t))$ do not differ significantly and can be considered to be approximately the same. The quality of this approximation improves when each update

$$\theta^{(k)}(i+1) = \theta^{(k)}(i) - \alpha \nabla \theta^{(k)}(i)$$

creates only a small displacement in the weight space. This can be controlled by suitably tuning the learning rate α . Qualitatively, the learning rate should decrease as the staleness in the system increases in order to reduce the divergence across the weights seen by the learners. The learning rate modulation of equation (6) achieves precisely this effect.

These results help define a principled approach for distributed training of neural networks: *the mini-batch size per learner should be reduced as more learners are added to the system in way that keeps $\mu\lambda$ product constant. In addition, the learning rate should be modulated to account for stale gradients.* In Table II, 1-softsync ($\sigma = 1$) protocol invariably shows the smallest training time for any $\mu\lambda$. This is expected,

σ	μ	λ	Synchronization protocol	Test error	Training time(s)
1	4	30	1-softsync	18.09%	1573
0	8	30	Hardsync	18.56%	1995
30	4	30	30-softsync	18.41%	2073
0	4	30	Hardsync	18.15%	2235
18	8	18	18-softsync	18.92%	2488

TABLE III
Results on CIFAR10 benchmark: TOP-5 BEST (σ, μ, λ) CONFIGURATIONS THAT ACHIEVE A COMBINATION OF LOW TEST ERROR AND SMALL TRAINING TIME.

since the 1-softsync protocol minimizes the traffic at the parameter server. Table II also shows that the test error increases monotonically with the $\mu\lambda$ product. These results reveal the scalability limits under the constraints of preserving the model accuracy. Since the smallest possible mini-batch size is 1, the maximum number of learners is bounded. This upper bound on the maximum number of learners can be relaxed if an inferior model is acceptable. Alternatively, it may be possible to reduce the test error for higher $\mu\lambda$ by running for more number of epochs. In such a scenario, adding more learners to the system may give diminishing improvements in the overall runtime. From machine learning perspective, this points to an interesting research direction on designing optimization algorithm and learning strategies that perform well with large mini-batch sizes.

D. Summary of results on CIFAR10 benchmark

Table III summarizes the results obtained on the CIFAR10 dataset using the Rudra-base system architecture. As a reference for comparison, the baseline configuration $(\sigma, \mu, \lambda) = (0, 128, 1)$ achieves a test error of 17.9% and takes 22,392 seconds to finish training 140 epochs.

E. Results on ImageNet benchmark

The large model size of the neural network used for the ImageNet benchmark and the associated computational cost of training this model prohibits an exhaustive state space exploration. The baseline configuration ($\mu = 256, \lambda = 1$) takes 54 hours/epoch. Guided by the results of section V-C, we first consider a configuration with $\mu = 16, \lambda = 18$ and employ the Rudra-base architecture with hardsync protocol (*base-hardsync*). This configuration performs training at the speed of ~ 330 minutes/epoch and achieves a top-5 error of 20.85%, matching the accuracy of the baseline configuration ($\mu = 256, \lambda = 1$, section IV-B).

The synchronization overheads associated with the hardsync protocol deteriorate the runtime performance and the training speed can be further improved by switching over to the 1-softsync protocol. Training using the 1-softsync protocol with mini-batch size of $\mu = 16$ and 18 learners takes ~ 270 minutes/epoch, reaching a top-1 (top-5) accuracy of 45.63% (22.08%) by the end of 30 epochs (*base-softsync*). For this particular benchmark, the training setup for the 1-softsync protocol differs from the hardsync protocol in certain subtle, but important ways. First, we use an adaptive learning rate

Configuration	Architecture	μ	λ	Synchronization protocol	Validation error(top-1)	Validation error (top-5)	Training time (minutes/epoch)
<i>base-hardsync</i>	Rudra-base	16	18	Hardsync	44.35%	20.85%	330
<i>base-softsync</i>	Rudra-base	16	18	1-softsync	45.63%	22.08%	270
<i>adv-softsync</i>	Rudra-adv	4	54	1-softsync	46.09%	22.44%	212
<i>adv*-softsync</i>	Rudra-adv*	4	54	1-softsync	46.53%	23.38%	125

TABLE IV

Results on ImageNet benchmark: VALIDATION ERROR AT THE END OF 30 EPOCHS AND TRAINING TIME PER EPOCH FOR DIFFERENT CONFIGURATIONS.

method (AdaGrad [8, 7]) to improve the stability of SGD when training using the 1-softsync protocol. Second, to improve convergence we adopt the strategy of warmstarting [21] the training procedure by initializing the network’s weights from a model trained with hardsync for 1 epoch.

Further improvement in the runtime performance may be obtained by adding more learners to the system. However, as observed in the previous section, increase in the number of learners needs to be accompanied by a reduction in the mini-batch size to prevent degradation in the accuracy of the trained model. The combination of a large number of learners and a small mini-batch size represents a scenario where the Rudra-base architecture may suffer from a bottleneck at the parameter server due to the frequent pushGradient and pullWeights requests. These effects are expected to be more pronounced for large models such as ImageNet. The Rudra-adv architecture alleviates these bottlenecks, to some extent, by implementing a parameter server group organized in a tree structure. $\lambda = 54$ learners, each processing a mini-batch size $\mu = 4$ trains at ~ 212 minutes/epoch when using Rudra-adv architecture and 1-softsync protocol (*adv-softsync*). As in the case of Rudra-base, the average staleness in the gradients is close to 1 and this configuration also achieves a top-1(top-5) error of 46.09%(22.44%).

The Rudra-adv* architecture improves the runtime further by preventing the computation at the learner from stalling on the parameter server. However, this improvement in performance comes at the cost of increasing the average staleness in the gradients, which may deteriorate the quality of the trained model. The previous configuration runs at ~ 125 minutes/epoch, but suffers an increase in the top-1 validation error (46.53%) when using Rudra-adv* architecture (*adv*-softsync*). Table IV summarizes the results obtained for the 4 configurations discussed above. It is worth mentioning that the configuration $\mu = 8$, $\lambda = 54$ performs significantly worse, producing a model that gives top-1 error of $> 50\%$ but trains at a speed of ~ 96 minutes/epoch. This supports our observation that scaling out to large number of learners must be accompanied by reducing the mini-batch size per learner so the quality of the trained model can be preserved.

Figure 8 compares the evolution of the top-1 validation error during training for the 4 different configuration summarized in Table IV. The training speed follows the order *adv*-softsync* $>$ *adv-softsync* $>$ *base-softsync* $>$ *base-hardsync*. As a result, *adv*-softsync* is the first configuration to hit the 48% validation error mark. Configurations other than *base-hardsync* show




Fig. 8. Results on ImageNet benchmark: Error on the validation set as a function of training time for the configurations listed in Table IV

marginally higher validation error compared with the baseline. As mentioned earlier, the experiments with 1-softsync protocol use AdaGrad to achieve stable convergence. It is well-documented [24, 21] that AdaGrad is sensitive to the initial setting on the learning rates. We speculate that tuning the initial learning rate can help recover the slight loss of accuracy for the 1-softsync runs.

VI. RELATED WORKS

To accelerate training of deep neural networks and handle large dataset and model size, many researchers have adopted GPU-based solutions, either single-GPU [16] or multi-GPU [26] GPUs provide enormous computing power and are particularly suited for the matrix multiplications which are the core of many deep learning implementations. However, the relatively limited memory available on GPUs may restrict their applicability to large model sizes.

Distbelief [7] pioneered scale-out deep learning on CPUs. Distbelief is built on tens of thousands of commodity PCs and employs both model parallelism via dividing a single model between learners, and data parallelism via model replication. To reduce traffic to the parameter server, Distbelief shards parameters over a parameter server group. Learners asynchronously push gradients and pull weights from the parameter server. The frequency of communication can be tuned via *npush* and *nfeteh* parameters.

More recently, Adam [4] employs a similar system architecture to DistBelief, while improving on Distbelief in two respects: (1) better system tuning, e.g. customized concurrent memory allocator, better linear algebra library implementation, and passing activation and error gradient vector instead of the weights update; and (2) leveraging the recent improvement in machine learning, in particular convolutional neural network to achieve better model accuracy.

In any parameter server based deep learning system [12], staleness will negatively impact model accuracy. Orthogonal to the system design, many researchers have proposed solutions to counter staleness in the system, such as bounding the staleness in the system [6] or changing optimization objective function, such as elastic averaging SGD [25]. In this paper, we empirically study how staleness affects the model accuracy and discover the heuristics that smaller mini-batch size can effectively counter system staleness. In our experiments, we derive this heuristics from a small problem size (e.g., CIFAR10) and we find this heuristic is applicable even to much larger problem size (e.g., ImageNet). Our finding coincides with a very recent theoretical paper [19], in which the authors prove that in order to achieve linear speedup using asynchronous protocol while maintaining good model accuracy, one needs to increase the number of weight updates conducted at the parameter server. In our system, this theoretical finding is equivalent to keeping constant number of training epochs while reducing the mini-batch size. To the best of our knowledge, our work is the first systematic study of the tradeoff between model accuracy and runtime performance for distributed deep learning.

VII. CONCLUSION

In this paper, we empirically studied the interplay of hyperparameter tuning and scale-out in three protocols for communicating model weights in asynchronous stochastic gradient descent. We divide the learning rate by the average staleness of gradients, resulting in faster convergence and lower test error. Our experiments show that the 1-softsync protocol (in which the parameter server accumulates λ gradients before updating the weights) minimizes gradient staleness and achieves the lowest runtime for a given test error. We found that to maintain a model accuracy, it is necessary to reduce the mini-batch size as the number of learners is increased. This suggests an upper limit on the level of parallelism that can be exploited for a given model, and consequently a need for algorithms that permit training over larger batch sizes.

ACKNOWLEDGEMENT

The work of Fei Wang is partially supported by National Science Foundation under Grant Number IIS-1650723.

REFERENCES

- [1] Mpi 3.0 standard. www.mpi-forum.org/docs/mpi-3.0/mpi30-report.pdf, 2012.
- [2] Y. Bengio. Practical recommendations for gradient-based training of deep architectures. In *Neural Networks: Tricks of the Trade*, pages 437–478. Springer, 2012.
- [3] T. M. Breuel. The effects of hyperparameters on SGD training of neural networks. *arXiv:1508.02788*, 2015.
- [4] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman. Project Adam: Building an efficient and scalable deep learning training system. *OSDI'14*, pages 571–582, 2014.
- [5] A. Coates, B. Huval, T. Wang, D. Wu, B. Catanzaro, and N. Andrew. Deep learning with cots hpc systems. In *Proceedings of the 30th ICML*, pages 1337–1345, 2013.
- [6] H. e. a. Cui. Exploiting bounded staleness to speed up big data analytics. In *USENIX ATC'14*, pages 37–48, 2014.
- [7] J. Dean, G. S. Corrado, R. Monga, K. Chen, M. Devin, Q. V. Le, M. Z. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Y. Ng. Large scale distributed deep networks. In *NIPS*, 2012.
- [8] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [9] O. R. et al. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, pages 1–42, 2015.
- [10] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, pages 249–256, 2010.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *arXiv preprint arXiv:1502.01852*, 2015.
- [12] Q. Ho, J. Cipar, H. Cui, S. Lee, J. K. Kim, P. B. Gibbons, G. A. Gibson, G. Ganger, and E. P. Xing. More effective distributed ML via a stale synchronous parallel parameter server. In *NIPS* 26, pages 1223–1231. 2013.
- [13] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [14] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [15] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Computer Science Department, University of Toronto, Tech. Rep.*, 1(4):7, 2009.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [17] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.
- [18] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [19] X. Lian, Y. Huang, Y. Li, and J. Liu. Asynchronous Parallel Stochastic Gradient for Nonconvex Optimization. *ArXiv e-prints*, June 2015.
- [20] M. Raynal and M. Singhal. Logical time: Capturing causality in distributed systems. *Computer*, 29(2):49–56, 1996.
- [21] A. Senior, G. Heigold, M. Ranzato, and K. Yang. An empirical study of learning rates in deep neural networks for speech recognition. In *ICASSP*, pages 6724–6728. IEEE, 2013.
- [22] A. Smola and S. Narayanamurthy. An architecture for parallel topic models. *Proc. VLDB Endow.*, 3(1-2):703–710, 2010.
- [23] I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th ICML*, pages 1139–1147, 2013.
- [24] M. D. Zeiler. Adadelta: An adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- [25] S. Zhang, A. Choromanska, and Y. LeCun. Deep learning with elastic averaging SGD. *CoRR*, abs/1412.6651, 2014.
- [26] Y. Zou, X. Jin, Y. Li, Z. Guo, E. Wang, and B. Xiao. Mariana: Tencent deep learning platform and its applications. *Proc. VLDB Endow.*, 7(13):1772–1777, 2014.